CDF Initiative Initial Deployment Plan For CAF Monitoring With Zabbix

Gerald Guglielmo

(CD-doc-2656)

Abstract

As part of the CDF Initiative it has been determined that CDF CAF operations could benefit by the deployment of a monitoring system to allow quicker assessment of the health of the CAF systems. This document is a plan for initially deploying the Zabbix monitoring system for CDF CAF operations monitoring.

Introduction

The CDF Initiative was initially created to investigate CDF CAF operations and determine where targeted effort could be applied to improve the operations environment by reducing the overall load on operations personnel. Early on it was determined that one of the impediments to understanding where to apply targeted effort was a lack of monitoring and issue tracking that could provide historical data on the types and frequency of problems in the CAF systems. Therefore it was determined that one of the first steps toward longer term improvement in operational load would be deployment of a monitoring system.

Prior to the start of deployment, a goal was set to find an existing monitoring system that would allow a phased deployment for CAF operations, starting with one set of monitoring tasks. The aim was not to do an evaluation of many candidate monitoring systems and pick the best one, but instead look for one, preferably with availability of local expertise, that would be sufficient for CAF operations monitoring. The result of this process was a determination that the Zabbix monitoring system (http://www.zabbix.com) did meet the defined requirements.

The goal for the first phase of deployment of Zabbix, configuration and commissioning to allow CAF operations monitoring of a subset of systems, is June 1, 2008. The details of the items to monitor in this first phase of deployment are defined in Schedd table in *CDF Initiative Monitor Initial Deployment List* (CD-doc-2655). Later phases of the deployment will be scheduled after experience is gained with the previously completed phase, with the intended goal of completing all phases by the end of September, 2008.

System Requirements

The CDF CAF systems provide a distributed computing environment for users around the world. These systems are often used 24 hours a day, 7 days a week and are a key component in carrying out the experiment's mission. Because of the complex nature of systems that integrate components from many different projects, some locally supported and others not, providing a high level of operational performance while minimizing the load in support personnel is a challenge. In this type of environment, a monitoring system that can allow support personnel to quickly assess the health of the systems and help isolate the cause of problems, and to do this remotely when necessary, is extremely important.

The Zabbix monitoring system relies on a set of agents that run remotely on the hosts to be monitoring. The job of these agents is to collect monitoring information and send it on to the Zabbix server. The agents come with a set of

predefined monitoring tasks that can be performed, and with support for registering and running user defined tasks. User defined tasks can in general be any user program or script that can be called by the Zabbix agent process and write their data to standard output. Registering of user tasks is fairly easy and is done by editing the configuration file for the Zabbix agent. Note that registering a task with an agent is only a declaration, it is not a request to actually perfrom the task. Requests for performing tasks come from the Zabbix server and in general must match the list of registered tasks in the agent.

The Zabbix server is responsibly for discovering the list of monitoring tasks from the Zabbix database and requesting the appropriate Zabbix agents to perform those tasks and return the data. The server than takes that data and stores it in the database. The server does not handle user requests to see the data, that instead is handled via a web interface to the database. Zabbix primarily supports Postgres and Mysql as backend databases.

User interaction with a running Zabbix system is through the web interface, which in turn talks to the database to retrieve information. This interface provided by a set of PHP scripts and served by an Apache web server. The simplest configuration, and the one recommended, is to run the Zabbix server, database, and web server on the same node. During the evaluation phase, by far the largest load on the system was from Postgres processes.

The storage database capacity, and thus disk storage capacity and backup requirements, seem to be very modest for CAF operations. Based on the number of items to check, the estimated frequency of checks, and the retention period, the maximum data storage needs are about 3 GB. Allowing for a factor of two safety margin means worst case expectations of 6 GB. This estimate includes the Groupcaf worker node checks which may or may not be included in the final monitoring system. They represent about 84% of the 2353 items to monitor every 20 minutes.

Either Postgres or Mysql database engines should be adequate for this task, and some form of backup will be necessary. A Linux server machine, running SLF 4.x or SLF 5.x, will be needed. According to the *Zabbix User's Manual,* the hardware specifications for the completed monitoring system including Groupcaf worker nodes is on par with a Intel Dual Core 6400 with 4GB of memory. However, for the initial phases of the deployment there will be only a few hundred items being monitored and a temporary machine with lower capabilities would work.

It is expected that the mechanism to install the user monitor programs and Zabbix agent configuration files on all of the remote nodes would be based on the standard system administration tools used in the Computing Division for similar system administration tasks. All user programs, Zabbix configuration files, etc. will be placed in CVS, or appropriate repository, and code deployments will be based on uniquely tagged versions.

Deliverables

There a several components to the monitoring system that comprise the basic infrastructure. In addition to the basic infrastructure there are installation and configuration tasks that need to be completed to commission the base infrastructure before it can be configured for CAF monitoring. Finally there is the specific infrastructure and configuration for CAF monitoring. As stated in the introduction, the deployment of CAF specific monitoring will be done in phases. The phase will include all of the base infrastructure and the monitoring checks listed in the Schedd table in *CDF Initiative Monitor Initial Deployment List* (CD-doc-2655). Later phases of the deployment will be scheduled after experience is gained with the previously completed phase, with the intended goal of completing all phases by the end of September, 2008.

(A) Hardware Server Platform

(A.1) Physical hardware

A server quality node with sufficient memory, cpu and disk space for hosting the web interface, database engine, and Zabbix server is required. As stated above, based on information from the *Zabbix User's Manual,* the hardware specifications for the completed monitoring system including Groupcaf worker nodes is on par with a Intel Dual Core 6400 with 4GB of memory. However, for the initial phases of the deployment there will be only a few hundred items being monitored and a temporary machine with lower capabilities would work. It is assumed that a standard network throughput for a server class machine will be sufficient. Other peripherals as required to allow following of standard practices and procedures of the hardware and operating system support teams is assumed.

(A.2) Operating Systematically

Zabbix runs under Scientific Linux Fermi 4.x, with either a later version of PHP or modifications to one of the PHP scripts, and it is expected to work under Scientific Linux Fermi 5.x. Either of these options are acceptable and the decision of which version is left to the installation and support team for the operating system. The desire it to keep the installation as standard as possible and consistent with the preferred version of the operating system support team, with only the absolutely necessary changes to allow Zabbix to function. For a discussion of details on possible changes needed under Scientific Linux Fermi 4.2, see Guidance on Installation and *Configuration of a Zabbix Monitoring System for CDF Offline* (CD-doc-2659).

(B) Database

(B.1) Database Engine

Either a Postgres or Mysql database installation will be sufficient. The decision on which of these two database engine to use should be based on which one the support team is most comfortable supporting. From a technical perspective there is nothing known at this time that makes either one a better choice.

(B.2) Database Configuration

The database should be configured to listen on local ports only unless some external constraint makes this unfeasible. The Zabbix infrastructure talks to the database either through a web server, or through the Zabbix server daemon. In both cases these processes are expected to run on the same host as the database engine, and thus local host access only is sufficient. The database should be installed and run under a standard account with standard permissions for access that are consistent with the Fermilab Computer Security policies. There should not be need for direct user access to the database, and this can be limited to the team supporting the database. The database repository should be on a file system with sufficient resources to meet the expected size based on the data retention estimates. Currently this size, with a safety factor of 2 is 6 GB. There should be a mechanism to easily recover and recreate the database configuration in the event of a system failure and loss of the database.

(B.3) Database Backups

While the data stored in the database will not be absolutely critical to the running of the experiment, the CAF specific configurations would be time consuming to reproduce from scratch and could severely impact the ability of the CAF team to support operations if it had to be reproduced from scratch. Therefore a backup strategy will need to be put in place. The database support team will be responsible for proposing a scheme and schedule, and implementing the scheme once approved.

(C) Zabbix Server

(C.1) Server Build and Installation

The Zabbix server will be built and installed on the monitoring system host. Note that the actual build of the server executable could be performed on a separate system with a consistent operating system environment.

(C.2) Server Configuration

The Zabbix server configuration should need only minor customization. This configuration should be saved some where it can be easily retrieved in the event of a system failure which causes the

production configuration to be lost. The server should be configured to run under an appropriate account and be started automatically on the system.

(D) Apache Web Server

(D.1) Configuration

The Zabbix user interface requires a web server that supports execution of PHP scripts. The Apache web server is recommended and needs to be configured to allow external connections, add a place in the directory tree to hold the Zabbix PHP scripts, and allow execution of PHP scripts.

(E) Zabbix Web Interface

(E.1) Configuration

(E.1.1) Accounts

There are three types of accounts that are generally of use for monitoring. The first is the super administration account with full privileges for basic Zabbix administration. Second is an administrator account for configuring CAF specific monitoring like defining items to monitor and graphs to generate. The third is an account that can view things like graphs and events, but has no privileges to modify existing settings or create new ones. Zabbix comes with an existing super administration account. At the very least one CAF administrator and one CAF reader account would need to be defined. Whether any other accounts are necessary depends on the authentication mechanism that is supported in the final system. Restrictions on who can access the CAF administrator account, or accounts, will be determined by the final agreement on who supports the Zabbix installation.

(F) Agent Infrastructure

(F.1) Zabbix Agents

(F.1.1) Agent Installation

Zabbix comes with agent daemons for collecting data and forwarding on to the server. These agents will need to be installed on each remote node that is being monitored. There needs to be a mechanism to install the agent code on each remote node to be monitored.

(F.1.2) Agent Configurations

The agents should be configured to run under an appropriate account and be started automatically on the system. There needs to be a mechanism for automatically installing configuration files on each of the remote nodes.

(F.2) Monitoring Scripts

While Zabbix comes with a large set of predefined checks it can perform, there are CAF specific checks that require user written scripts that the Zabbix agent can execute. These scripts will be written by the CAF team with assistance from CD where appropriate and agreed. These scripts will be stored in a revision control system and deployed from uniquely tagged releases. For each of these scripts the Zabbix agent configuration will need to be modified, to register the action to perform, and those changes will be provided by the CAF team with assistance from CD. There will need to be a mechanism to deploy these scripts to each remote node to be monitored.

(G) First Release Phase Monitoring System

(G.1) CAF Team Accessible Monitoring System

An operational monitoring system that provides monitoring of items identified for phase 1 while be commissioned for viewable access by the CAF team and other members of the deployment team. This initial restricted period is to allow some shake out of the system before opening it up for wider use. Assuming the system is stable, opening up the system for wider viewing should be no longer than one week later. Initial milestone is beginning of June, 2008.

(G.2) Open Viewing Access

The monitoring system will be made open for viewing to CAF users following the initial commissioning phase. Target date for opening the system access is end of the first week of June, 2008. The CAF team with assistance from CD will watch the performance of the monitoring system to see if wider access causes any issues.

(H) Additional Release Phases of Monitoring System

(H.1) Open Viewing Access

Based on experience gained during the first release phase, it should not be necessary, and probably not practical, to have an internal only release. A suitable set of additional checks, preferably one or two new templates, will be added to the system and made available for user viewing. The system will then be monitored for performance issues and planning for the next cycle will happen in parallel.

Risks

Still need to write this section. Check back later... GMG